

AMENDMENTS TO THE CLAIMS:

1. (Previously presented) A method of improving at least one of speed and efficiency when executing a linear algebra subroutine on a computer having a memory hierarchical structure including at least one cache, said computer having M levels of caches and a main memory, said method comprising:

determining, based on sizes, for a level 3 matrix multiplication processing, which matrix will have data for a submatrix block residing in a lower level cache of said computer and which two matrices will have data for submatrix blocks residing in at least one higher level cache or a memory;

selecting, from a plurality of six kernels, two kernels optimal to use for executing said level 3 matrix multiplication processing as data streams from different levels of said M levels of cache, such that said processor will switch back and forth between said two selected kernels as steaming data traverses said different levels of cache and

streaming data from said selected two matrices, for executing said level 3 matrix multiplication processing, so that said submatrix block residing in said lower level cache remains resident in said lower level cache.

2. (Previously presented) The method of claim 1, wherein said lower level cache comprises an L1 cache and said higher level cache comprises an L2 cache.

3. (Previously presented) The method of claim 1, wherein said determining said matrix to be stored in said lower level cache comprises determining which of the three matrices has a smallest size.

4-5. (Canceled)

6. (Previously presented) The method of claim 2, wherein data for said second matrix and said third matrix streams into said L1 cache from said L2 cache such that said data from said second matrix and said third matrix streams in a vector format into said L1 cache.

7. (Previously presented) The method of claim 1, wherein said linear algebra subroutine comprises a substitute of a subroutine from LAPACK (Linear Algebra PACKage).

8. (Previously presented) The method of claim 7, wherein said substitute subroutine comprises a BLAS (Basic Linear Algebra Subroutine) Level 3 routine or a BLAS Level 3 kernel routine.

9. (Currently amended) An apparatus, comprising:

a memory system to store matrix data for a level 3 matrix multiplication processing using data from a first matrix, a second matrix, and a third matrix, said memory system including at least one cache; and

a processor to perform said level 3 matrix multiplication processing, wherein data from one of said first matrix, said second matrix, and said third matrix is stored as a submatrix block resident in a lower level cache in a matrix format and data from a

remaining two matrices is stored as submatrix blocks in said memory system at a level in said memory system higher than said lower level cache,

said processor preliminarily selecting, based on sizes, which matrix will have said submatrix block stored in said lower level cache and which said two matrices will have submatrix blocks stored in said higher level,

said data from said selected two matrices being streamed through said lower level cache into said processor, as required by said level 3 matrix multiplication processing, so that said submatrix block stored in said lower level cache remains resident in said lower level cache,

wherein said computer comprises M levels of caches and a main memory, said processor further preliminarily selecting, from a plurality of six kernels, two kernels optimal to use for executing said level 3 matrix multiplication processing as data streams from different levels of said M levels of cache, such that said processor switches back and forth between said two selected kernels as streaming data traverses said different levels of cache.

10. (Previously presented) The apparatus of claim 9, wherein said processor selects a smallest of said first, second, and third matrices to be said matrix to have data residing in said first level cache.

11. (Previously presented) The apparatus of claim 9, wherein said level 3 matrix multiplication comprises one or more subroutines substitute to a subroutines from LAPACK (Linear Algebra PACKage).

12. (Previously presented) The apparatus of claim 11, wherein said substitute subroutine comprises a BLAS (Basic Linear Algebra Subroutine) Level 3 routine or a BLAS Level 3 kernel routine.

13-23. (Canceled)